

**Author:** Claire Coleman

**Team:** Customer Service & Communication Team

**Date:** 7/03/08

**Version:** 1.0

### Introduction

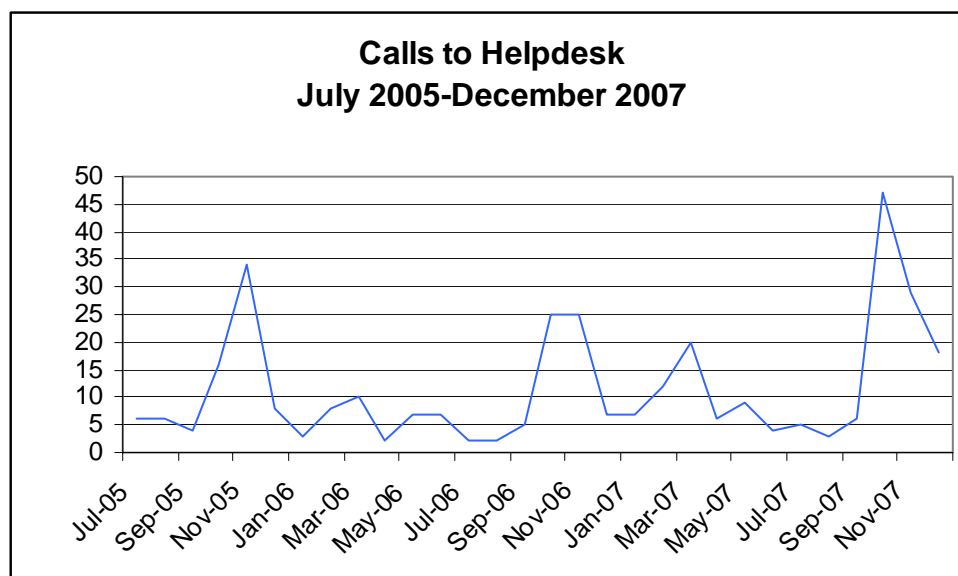
The University of Sheffield's High Performance Computing (HPC) facility is provided by CiCS. It consists of both hardware and software (support and provision) to researchers at the University. The CiCS Sheffield HPC server is called Iceberg.

Iceberg forms the Sheffield node of the White Rose Grid (WRG). The White Rose Grid is a large research computing infrastructure comprised of grid computing facilities from the Universities of Leeds, York and Sheffield.

This service review was carried out for two reasons. Firstly there is a need to consult with HPC users to ensure that ongoing developments are appropriate. Secondly, due to other factors, there has been more downtime in service in the last year and we were interested on the effects this had had on users.

A questionnaire was developed with the CiCS Research Computing Group to send to people who had access to the HPC Service.

### Helpdesk Statistics





The number of calls to the helpdesk about the HPC service show peaks in October and November when new students start. It also shows a growth in those calls over the last 3 years.

The number of calls are relatively small compared to the number of people with access to the service (764) suggesting that either the service is reasonably easy to use and/or that the users are already very skilled.

### Method

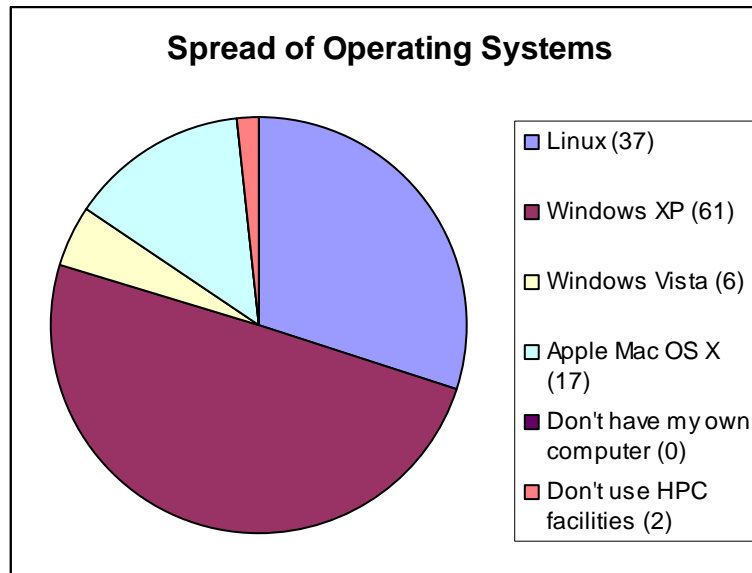
A questionnaire was emailed to 738 people who were registered to use the HPC service. It asked specific questions (see full results appendix 1) and gave free text boxes for users comments (see appendix 2).

### Results

The user population was classified as follows:

Classification of Users	Numbers sent Questionnaire	Number responding to Questionnaire	Percentage of user responses
Staff	283	45	16%
Post-Graduate Research	203	30	15%
Post-Graduate Taught	43	2	5%
Under-Graduate	121	7	6%
Other	88	2	2%
<b>Total</b>	<b>738</b>	<b>86</b>	<b>N/A</b>

The percentage of responses from each user group indicates that the HPC Service is of higher importance to Research Post Graduate Students and Staff.



The majority of users used Windows XP (70.9%) although 43% used Linux.

The majority used the service weekly although there was a fairly even spread between daily, weekly, monthly and rarely as responses.

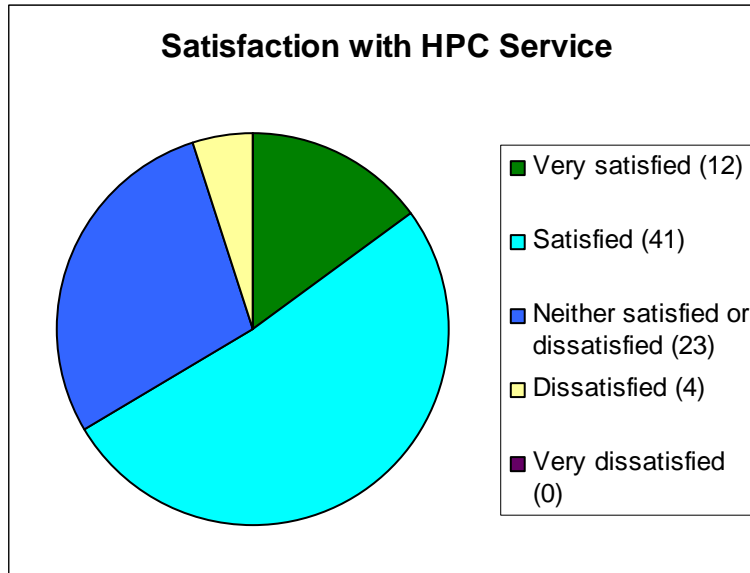
78.8% use the HPC service from their office on campus but a significant number use it from home (52.5%).

45% saw the HPC service as vital to their work.

66% use CiCS provided software and 52.5% write their own applications.

The majority (55%) use the service for both batch and interactive jobs.

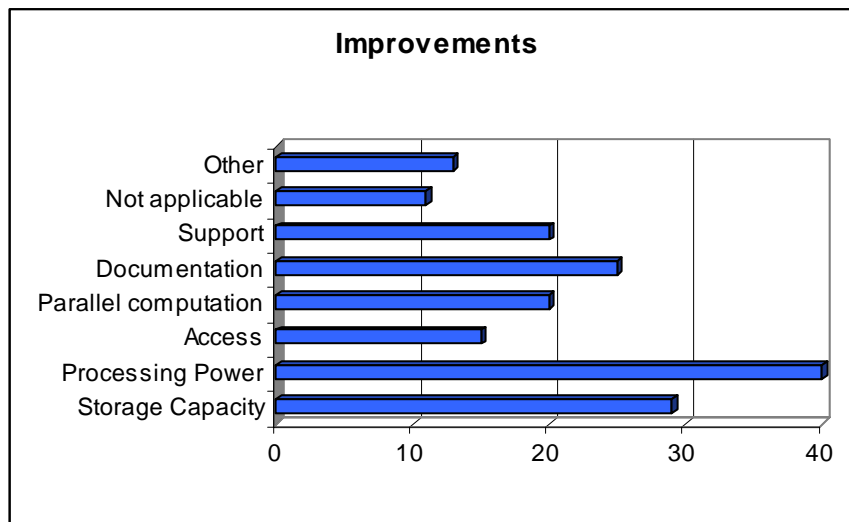
There was an awareness of other HPC services (91.4% White Rose Grid, 50.7% the National Grid service, 30% HPCX) but 46.3% did not see using another service as applicable to their needs.



There were only 4 people dissatisfied and no-one was very dissatisfied.

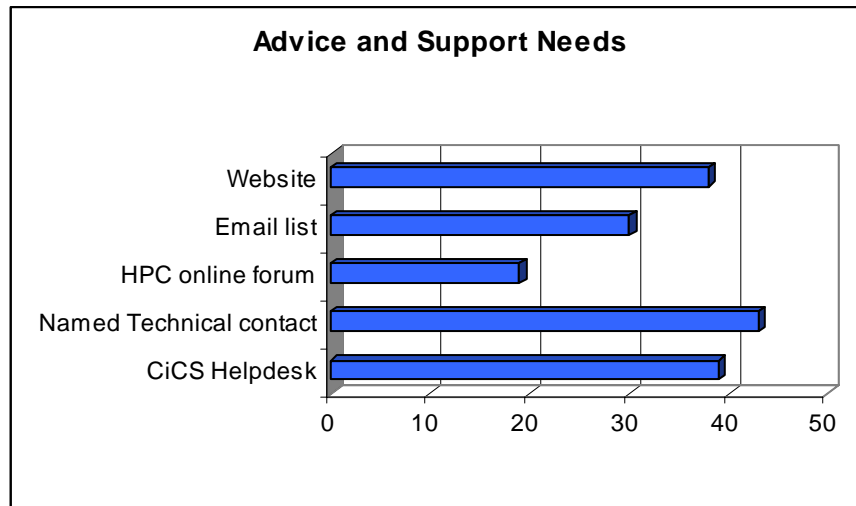

90.9% were happy with the amount of downtime and 84.9% were satisfied with communication about downtime. 89.7% want to be informed of downtime by email.

The table below shows improvements that users selected.



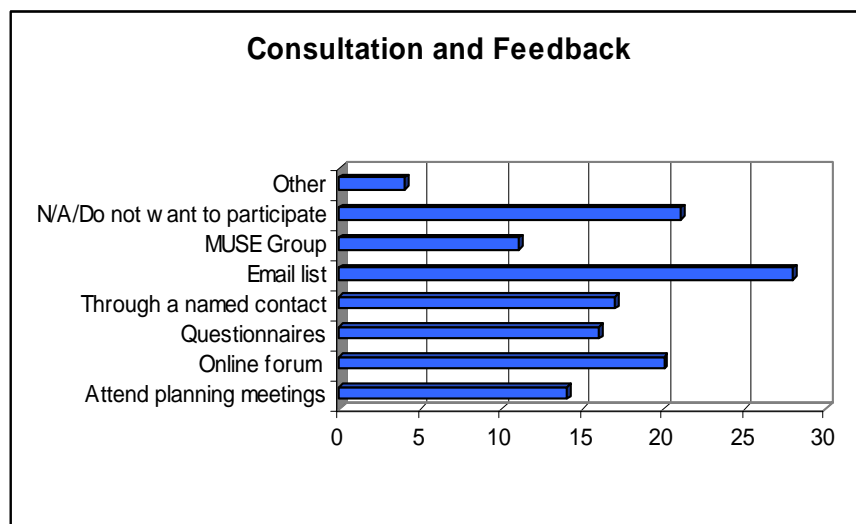
The main improvements people want are in processing power and storage capacity.

The table below shows what users want in terms of advice and support.



There is a fairly even spread between a named technical contact, helpdesk and the website.

The table below shows how users want to communicate their views.



Users were fairly evenly split between a number of methods.



## Summary

The HPC Service is well regarded. 66% of HPC users are satisfied or very satisfied with the HPC service CiCS provides. 90.0% were happy with the amount of downtime and 84.9% happy with the way in which downtime was communicated.

The satisfaction with downtime and communication is good because there was a larger amount of downtime in the last year caused by generator installation work and power cuts. There should be less downtime next year.

There are suggestions for improvement particularly in processing power, storage and documentation. There is an ongoing programme of service improvement at the moment.

There are a significant number of people using the service from home so documentation supporting this should be clear and accessible. New users also require more in depth support/documentation. This is reflected in the peaks in calls to helpdesk at the start of the academic year.

For communication it seems clear that information about downtime should be sent via email and support given in three ways – helpdesk, named support person and the website.

There was a fairly even split covering how users want to be involved in ongoing developments suggested offering a blend of activities would be best.


## Recommendations

- 1) Look into raising processing power
- 2) Expand storage capacity
- 3) Expand and improve documentation provided
- 4) Examine queue/priority management
- 5) Look into guidance and support for new users
- 6) Develop programme of consultation and feedback to feed into ongoing planning.

## Improvements already planned/in progress

- 1) Look into raising processing power

CICS allocates £75K per annum for updating existing research computing facilities. This has resulted in the replacement of the SUN V20 boxes with SunX2200 boxes and 2 dual CPU AMD opteron processors.



The latest procurement resulted in the provision of 32 SunX2200 nodes with 16GB of memory and two dual core AMD opteron (Rev. F) processors rated at 2.6GHz

A second tranche will result in the purchase of a further 64 X2200 nodes. Bring the total CPU power to 384 cores. The existing 19 V40 boxes have 4 AMD opteron single core CPU's and a myrinet interconnect, these are for dedicated parallel processing.

The system operating system has been updated to Linux 5.1 a future upgrade (September 2008) will result in the replacement of the current V40's.

Additional resource may be provided through the purchase of nodes for dedicated usage by research groups. Income generated through this method is used to procure additional compute resource.

## 2) Expand storage capacity

An extra 5.5TB storage has been ZFS mounted using the Sun 3510/3511 RAID storage system. This resulted in providing users with, initially, 5GB of backed-up storage allocation in their home area. A substantially increased quota of 50GB per user was provided for the data area.

The RCG is investigating future storage systems.

## 3) Expand and improve documentation provided

A new web site was created in summer 2007 work to keep this up to date and to provide further content with improved user documentation is in progress.

## 4) Examine queue/priority management


The RCG monitor the status of the cluster and its queues the queue and scheduler configurations are monitored and reconfigured to enable fair and optimal use of the compute cluster. Work is also undertaken to review:

- The sun grid engine fair sharing mechanism
- Other schedulers e.g. torque produced by moab

## 5) Look into guidance and support for new users

The RCG is planning to increase availability of training sessions and material by making modules available through MOLE (my online learning environment).

- Introduction to LINUX and HPC
- Introduction to Matlab



Current training modules and hands on training sessions are provided through the research training programme additional training sessions are provided subject to requests made through the learning management system and through an application form available on the training section of the web site.

- 6) Develop programme of consultation and feedback to feed into ongoing planning.

The RCG works with the CICS Customer Service and Communication team to liaise with the user community. The following mechanisms are available:

- Communicate requirements to the RCG through central help desk
- E-mail [iceberg-admins@lists.shef.ac.uk](mailto:iceberg-admins@lists.shef.ac.uk)
- News items are available by typing news when a user logs into iceberg
- The updated news is available on the web site at <http://www.shef.ac.uk/wrgrid/iceberg/news.html>
- All users are contacted using [iceberg-announce@lists.shef.ac.uk](mailto:iceberg-announce@lists.shef.ac.uk)

CICS have recently established a research computing advisory group which comprises a panel of leading research groups that utilise the compute cluster.

The group is considers the following issues:

- Procurement
- Cluster Configuration
- Applications
- Training and Support needs
- e-Science infrastructure requirements

- 7) Provision of additional applications

The following applications have been requested for updating or installation and are under investigation or have been installed:

- OpenBUGS
- Parallel fluent
- MPICH over Ethernet to enable parallel computing on the new workers
- Matlab distributed computing engine



## Appendices

**Appendix 1** – see HPCAppendix1.pdf

## Appendix 2

**Choose the category that best describes you. Other (please specify)**

phd

Retired Academic Staff

**What Operating System do you use to access the HPC Facility? Other (please specify)**

Cygwin under Windows XP

**Which Facility do you use? Other (please specify)**

matlab on iceberg

SSH Secure Shell Client

Major international scientific codes such as DLPOLY, CASTEP

**Where do you use the HPC Facilities? Other (please specify)**

From Leeds University

Remote research facilities and on the road

Shared office on campus

Use facilities for Summer School for 2 weeks annually

**Could you carry out your work without the HPC Facility? Other (please specify)**

Could do own research without HPC facilities. Could not run the Summer School without them

I have occasional jobs that I couldn't run otherwise

Project/task dependent. we've not had much need currently but we may have more soon

With effort, could now run my programs on my pc. Basic programs developed 15 or so years ago when pc's no so advanced.

With some difficulty. I cannot afford software such as MATLAB myself.

Without Sheffield's yes - provided we had good local provision at Leeds

Would be a lot harder

Yes - but more slowly

Yes - I have access to other HPC machines and I could apply for my own cluster, but being able to use local resources makes it all \*much\* easier.

Yes but results would have to be scaled down due to limited resources.

Yes, but I would need a better PC and I would still be limited in the amount of work I could do.

Yes, but it helps a lot! (could only solve much smaller problems otherwise)

Yes, but it would slow my research (and especially that of my PhD students) down a lot.

Yes, but some of it would be a lot slower

Yes, but would be hard

**Are you aware of other HPC Facilities? Other (please specify)**

everest

HECTOR

UKMHD Grid



The  
University  
Of  
Sheffield.

## Corporate Information & Computing Services

High Performance Computing Report  
March 2008

### **If you used other HPC facilities which would you use? Other (please specify)**

All  
don't know  
don't know  
HECTOR  
I don't have enough knowledge to make a choice  
It would depend on the computing problem  
Own (research group) cluster  
see above  
special Diamond or Daresbury software  
White Rose Grid, but may consider others if I knew more about them.

### **What do you feel could be improved? Other (please specify)**

BUGS doesn't work!

Communication. Specifically, the information provided in real-time when there is a system problem or a problem with my account (my account was suspended when I was over quota, and I didn't even get an email to let me know, it just stopped giving me access). On other occasions the system has been down for several hours in office time with no explanation. It is obviously annoying for the admins to have me email or ring them to find out if there is a general problem or a problem I have caused. This could be saved if there was a more general email sent out quickly if the system crashes (or a web update, as some people would regard such emails as spam).

Fortran compiler with easy support for quad precision arithmetic  
I would like to use it, but no-one has ever informed me that I could.  
Increased reliability.

Libraries and software could do with being kept more upto date.  
Network acces for huge (gigabyte) file transfer.

reliability

Reliability - down time on iceberg

Reliability!

Scheduling

Search engine (web-based) for interrogating library-based resources to be made available via web connection e.g. X-Windows Also see 12 below

Would be nice to be able to have more jobs running when the ICEBERG is quiet. Often processors seem to be idle, and yet jobs are still queued to be executed later


### **Please tell us more about any improvements you want**

As a semi-technical user I have welcomed mike's support of previous projects to make Grid technology accessible to us. It remains clear that more development is needed in this area. However, we recognise the scale of the task and the appropriate funding needed in this area.

Better communication to people to don't use it but might find it useful?

Faster processing. Stable connection to programs e.g. LS-Dyna (impossible to access from some machines, difficult from many). More rapid support (my last query took ver two months to get a response, by which time I had acquired my own licence to run locally).

Have experienced many problems with iceberg which I never experienced with previous sun systems (titania etc). I suspect that the Linux OS requires more intervention from the systems administrator than the sun systems did (e.g. clearing cache etc). Need capable staff!



The  
University  
Of  
Sheffield.

## Corporate Information & Computing Services

High Performance Computing Report  
March 2008

I want to be able to use backspace again (rather than ctrl-H) to delete previous characters when using a Qsh interactive window!

I wanted to run some analysis in the statistical language R (<http://www.r-project.org/>) using the batch submission and not having to login using an X session and leave that open. There is an R packages called GDD which allows this (see CRAN archives at <http://cran.r-project.org/>). However, R was out of date (at the time R-2.3 was installed and current version was R-2.5.0). The analysis packages (haplo.stats() again see CRAN) required  $\geq$ R-2.4.1. This was resolved by installing R locally, however, GDD installation locally required gdlb-config and the version installed was 2.0.28 whilst which didn't work (and was outdated since the current version was 2.0.35 and this worked on three systems that I used it on).

I would like to see an administrator take over the installation and maintenance of R.

It would be good if you could provide better basic guidance for students who have no concept of using a command line based system. A lot of the younger ones have never used anything except Windows-based GUIs. A list of basic commands would be helpful.

It would be good to increase the disk quota available - but I suspect that users always want more space! It is sometimes difficult to get MSc and PhD students started with HPC resources, and more online documentation would be a big help.

make it more accessible with GUI.

More CPUs and better interconnects would be nice. Compilers, access and how-to-use-Iceberg docs and courses are all excellent.

More, faster chips. I and my students tend to run lots of single processor jobs, so fast interconnect etc. isn't important.

No quota limits on non-backed up storage. Access to local storage on nodes.

Open-Ended Response

Our group (the ABRG) makes use of its own cluster. This uses the matlab DCE and a neural simulation engine GENESIS.

Please get BUGS working

Suggestions I have made regarding the service in the past have been acted on reasonably quickly (which is nice), so I will continue to let the admin team know if there are any large problems as and when they occur.

The number of queued processes must have a boundary per person. I mean that if somebody put 300 processes 20 will begin to run and the other 280 will appear in the queue. So the queue is becoming very large. But if you a reasonable number of queued per person the queue will be more readable

Tweaking either the priority algorithm (so it doesn't run x number of jobs by one person at once, for instance) or changing the job submission policy. Either way, they're relatively minor points. My main interest would be to increase the processing power available.

virtualized Multi-core nodes


When I first started using the facilities I found it a bit daunting and confusing, that was over a year ago. It's difficult to remember what my exact problems were but maybe more beginner friendly documentation would have been helpful.

workstations strategically placed within main buildings. Currently, groups must supply such equipment form own funds

### **What problems were there with system availability and how did they affect you?**

Had to wait a week to queue up a job of mine, no biggie really.

have not used system very much to date - this may change in the next months



The  
University  
Of  
Sheffield.

## Corporate Information & Computing Services

High Performance Computing Report  
March 2008

I was aware of planned downtime and the recent migration of the scratch area, and I was able to plan my work around them.

IP links / log-in sessions sometimes drop out unexpectedly - this might be a university network problem though

large jobs are far too long in queue

License server often unavailable, which prevents running applications such as matlab and pgi compilers.

Need more guidance and support for beginning users

None

None that really affected me.

Occasionally, jobs would be queued for a long time. I found this was fairly rare.

Open-Ended Response

Programmes stopping sometimes and the queue sometimes were longer ...

relatively small disk quotas, and lot of downtime

see above.

Students have reported downtime on iceberg and sometimes poor access to certain queues

There has been a few time when I have been told that connection hasn't been possible (wouldn't schedule) and if I could try again at a later time.

There seems to have been a large amount of downtime.

There was a close clash with the Summer School which would have been disasterous if the work was not shifted.

There was little downtime and bearable queue time so I do not believe it to have affected me.

### **How do you prefer to be informed about system downtime? Other (please specify)**

MUSE headlines

Not applicable - do not intend to use the service

RSS feed

warning on screen 2-5 minutes before shutdown

### **How would you like to participate in the ongoing development of the HPC facilities? Other (please specify)**

a competent focus group/steering cttee needs to be set up from major stakeholders

I would be happy to attend planning meetings if they were relevant

Much of the above, depending on what is appropriate for me to respond to.

WOuld participate, but will be leaving soon, so there seems little point

### **Thank you for filling in the HPC service review questionnaire. We will put the results of this questionnaire on our website. Please tell us anything else you want to feedback about the HPC facilities.**


Good job guys please keep it up.

i appreciate ur efforts in running iceberg

I found the CICS/ICOSS taverna workflow workshop very helpful- I think more examples like this across the university are important in building the community. Cheers, Robin :)

Open-Ended Response

Overall I am very happy with the current levels of provision, and timely upgrades. I would like to see this continue to be freely available. I think that once users have to pay then it becomes very tempting to obtain your own system. This approach has advantages, but there is a huge sysadmin cost.



The  
University  
Of  
Sheffield.

## Corporate Information & Computing Services

High Performance Computing Report  
March 2008

The recent increase in storage space is very welcome; I feel that the HPC team should now focus on giving Iceberg more and faster CPUs, and improving the interconnect for MPI. Impressed with what I've used so far though.

well done for the good work that you do.